PM7322

RCMP
IN
EGRESS OPERATION

Issue 2: March 13, 1996

TABLE OF CONTENTS

ABBREVIATIONS

ABR	Available Bit Rate service
CDV	Cell Delay Variation
NNI	Network Network Interface
OAM	Operations, Administrative and Maintenance
PDH	Pleisiosynchronous Digital Hierarchy
PHY	PHYsical Layer
QoS	Quality of Service
RCMP	Routing, Control, Monitoring and Policing standard product. *Note: In this document, the term RCMP is used to refer to both the RCMP-800 (PM7322) and the RCMP-200 (PM7323) devices*
RCMP-800	PM7322 RCMP-800 standard product, which can process up to 800Mbps data. It has a 16-bit SCI-PHY interface running at a maximum rate of 50MHz.
RCMP-200	PM7323 RCMP-200 standard product, which can process up to 200Mbps data. It has a 8-bit SCI-PHY interface running at a maximum rate of 25MHz.
RM	Resource Management
SCI-PHY	SATURN-Compliant Interface for ATM PHY Layer. Please refer to the document, PMC-940102, "SATURN Compliant Interface for ATM PHY Layer and ATM Layer Devices, Level 2", October 1995.
UNI	User Network Interface
VC	Virtual Channel

REFERENCE

- ITU-T Recommendation I.371, "Traffic Control and Congestion Control in B-ISDN", March 1993.
- ATM Forum "Traffic Management Specification V4.0", December 1995.
- PMC-940102, "SATURN Compliant Interface for ATM PHY Layer and ATM Layer Devices, Level 2", October 1995.
- PMC-940903, "ATM Layer Routing Control, Monitoring and Policing Standard Product Engineering Document", August 1995.
- PMC-941027, "S/UNI-622 Databook".
- PMC-931110, "S/UNI-155-LITE Databook".
- PMC-950449, "S/UNI-MPH Databook".

INTRODUCTION

Although the RCMP is intended primarily for the ingress direction application in switches, it can also be used in the egress direction to provide address translation, OAM processing, policing and monitoring functions.

A typical egress application is shown in Fig. 1. ATM traffic comes from the switch core into the RCMP, which can translate the switch core specific cell header into one that conforms to either the UNI or NNI specification. Extended cell format[1] could have been used in the switch core for routing purposes; the RCMP can remove the appended cell octets to produce 53-octet ATM cells for the PHY layer. Also, the RCMP can perform multicasting; the advantage of doing multicast in egress is that it avoids congesting the switch. OAM cells can be extracted, generated or inserted. An example is sending performance monitorting cells on the multicast VC's. The RCMP can perform policing to ensure rate conformance in each VC after the switching operation. Finally, the RCMP can gather performance monitoring data, such as counts for high/low priority cells, non-compliant cells, errored cells. An application of the policing and performance monitoring function is at the NNI between networks such that the traffic conformance can be checked for diagnostics and fault isolation purposes.

At the output of the RCMP, a traffic shaper can be used to alter the cell traffic characteristics on either a per-VC or per-PHY basis, in order for the network provider to meet the Quality-of-service (QoS) at the output of the switching network, or for the user to conform to the Traffic Contract. Examples of traffic shaping functions are peak cell rate reduction, burst length limiting, reduction of cell-delay variation (CDV) by suitably spacing cells in time. Note that traffic-shaping functions are optional[2]. Cell traffic from the shaper is then processed by the PHY device to map cells into the physical carrier format.

Fig. 1 Use of RCMP in the Egress Direction

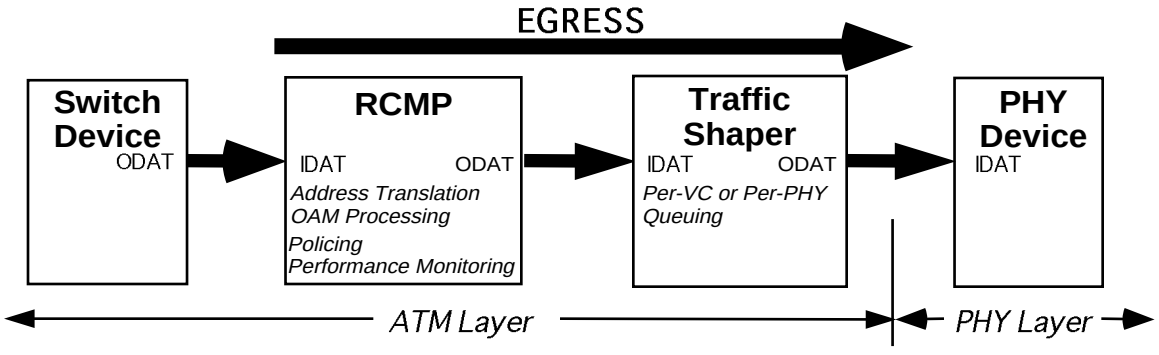

[1]Refer to PMC-940102, "SATURN Compliant Interface for ATM PHY Layer and ATM Layer Devices, Level 2", October 1995.

[2]Refer to ITU-T I.371, and ATM Forum Traffic Management Specification Version 4.0.

The purpose of this application note is to describe the various interfaces between the switch core and the RCMP, and between the RCMP and the downstream device, which can be either a traffic shaper or a PHY device.

INTERFACE DEFINITIONS

The RCMP uses the SCI-PHY Level 2 interface[3] at its input and output. There are two types of interfaces, depending on the direction of the read-enable signal. If the interface sources the read-enable signal, it is referred to as the *master*. If the interface receives the read-enable signal, it is referred to as the *slave*. Note that this is independent of whether the data (the cell traffic) is being sent or received.

Normally, a master interface is directly connected to a slave. However, this is not always the case, especially with the many existing proprietary interface designs, such as those found in the switch core. Therefore, it is necessary to understand how master-master and slave-slave interfaces work. The following two sections describe in detail the logic and timing considerations for these two interfaces.

MASTER-MASTER INTERFACE

A master interface can be connected to another master interface using a synchronous FIFO, as shown in Fig. 2, which is an example of connecting the RCMP[4] input interface to the output of the switch core. The timing diagram is shown in Fig. 3.

The FIFO acts as a buffer that can store and forward cells. This allows both master interfaces to independently control when to write out and read in a cell using the read- and write- enable signals respectively. The FIFO conveys the availability of data through the PAFB and PAEB signals, which are Programmable active-low Almost-Full and Almost-Empty signals. These signals control the OCA and ICA signals. On the switch side, PAFB is programmed to be asserted if the FIFO contains 38-64 words. With the FIFO being 64-word deep, this means that OCA is asserted (high) only if the FIFO has space for at least 27 words. Similarly on the RCMP side, the ICA (PAEB) signal is only asserted if the FIFO contains at least 27 words (ie. an entire cell is available in the FIFO). To support the extended cell format for SCI-PHY Level 2, all that is needed is programming the thresholds for the PAFB and PAEB signals to accomodate the extra words in a cell.

The two clocks on the switch and RCMP interface can be asynchronous. However, in most applications, these clocks are from the same source.

[3] This is backwards compatible with the UTOPIA interface, with the main difference being the support for the extended cell format. Refer to PMC-940102

[4] A 16-bit interface is shown in the figures, which applies to RCMP-800 only. However, these interfaces also apply to RCMP-200, with the difference being that the RCMP-200 has a 8-bit interface running at 25MHz.

Fig. 2 Master-master interface Circuit

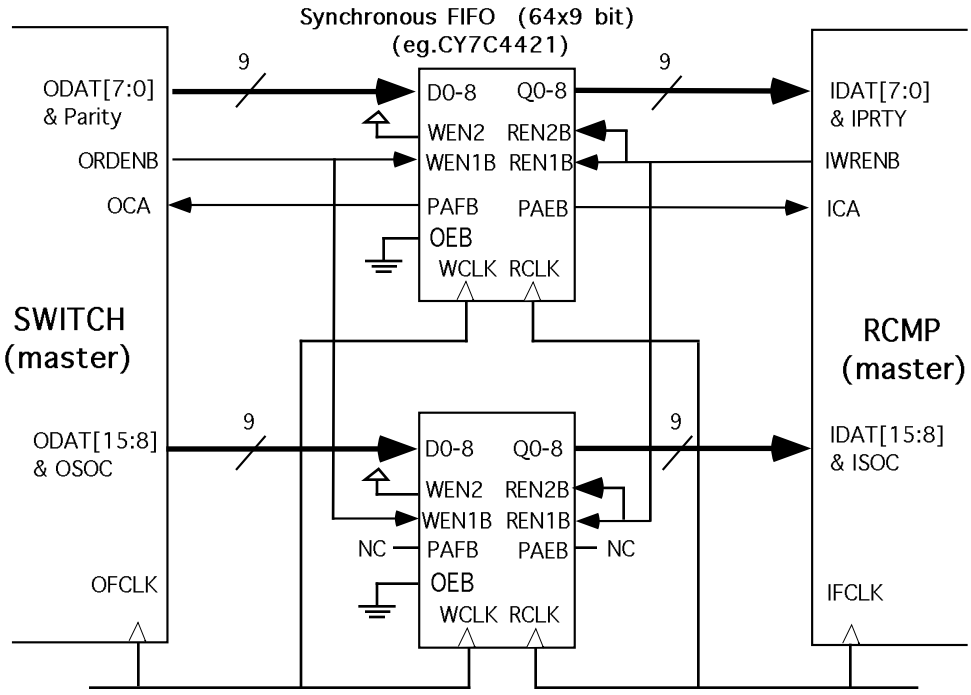

Fig. 3 Master-master interface timing

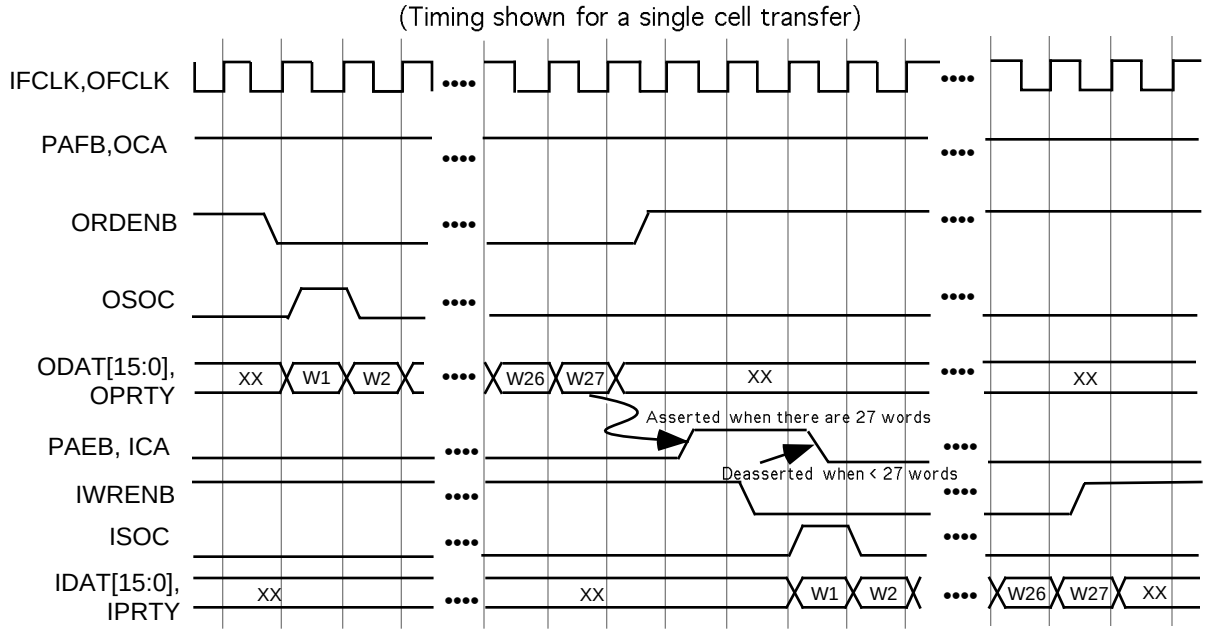

In this example, the 16-bit data bus with a single parity bit and the SOC signal are partitioned nicely into two 64x9 bit FIFO's. Other cases are less convenient, such as 1) 8-bit data bus with parity and SOC, making a total of 10 bits, and 2) using two parity bits instead of one in the 16-bit interface. Because synchronous FIFO's only come in either 9 or 18 bit sizes, there is no choice but to use an extra FIFO to accomodate the extra bit signals.

The RCMP can be interfaced directly to 4 switch devices using this scheme. Each switch device will have its own dedicated FIFOs. The data output from the FIFO will all share a single bus into the RCMP. The output enables of these FIFO's will have to be controlled by the IWRENB signals from the RCMP to tristate the FIFO data outputs. Multi-PHY interface using address-polling is not supported by this scheme.

Cell-based vs Word-based handshake

It is assumed that both interfaces use the cell-based handshake[5], where the switch can only start a cell transfer when there is at least 27 words of space in the FIFO, and the RCMP can only start a cell transfer when there are at least 27 words available in the FIFO. In order to allow continuous cell transfer, the FIFO must be at least 2-cell deep, which is the case as shown in Fig. 2. The reason for this is due to the potential dead time between when a cell has been written into the FIFO and when enough words of the cell is read out to provide space (27 words) for a new cell to be written in. Therefore, if the FIFO has a depth of less than 2 cells, the dead time will appear since OCA will not be asserted while the FIFO is outputting a cell that has just been written in. Once enough words have been outputted such that there is space for 27 words, OCA will be asserted. Now, with a 2-cell FIFO, there will be no dead time; ie. OCA will always be asserted since there is always space for a cell[6].

If the switch uses the word-based handshake[7], then a 2-cell FIFO will not be required. In this case, the PAFB needs to be programmed to be asserted when the FIFO contains at least 1 word. This means that OCA is high whenever there is space for one word in the FIFO. The switch only sends out one word at a time. Therefore, if the RCMP can always read cells whenever they are available, the FIFO will always have space for the switch to continuously send cells.

SLAVE-SLAVE INTERFACE

The slave-slave interface is quite straight forward, as shown in Fig. 5. The timing diagram is shown in Fig. 6.

[5]Cell-based handshake is where a device will start a complete cell transfer once a cell-available signal (eg. OCA, ICA) is asserted. The device does not stop until the entire cell is transferred.
[6]This is assuming the RCMP is continuously reading out the cells that are available in the FIFO.
[7]Word-base handshake is where a device only send one word every time a "word"-available signal is asserted.

Fig. 5 Slave-slave interface ciruit

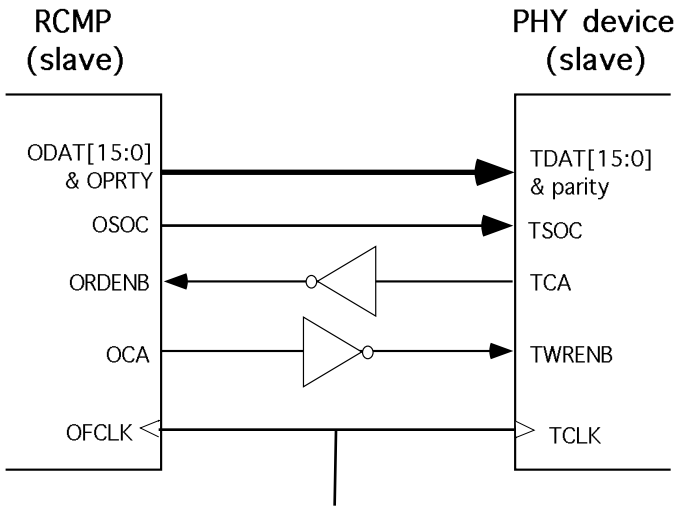

Fig. 6 Slave-slave interface timing

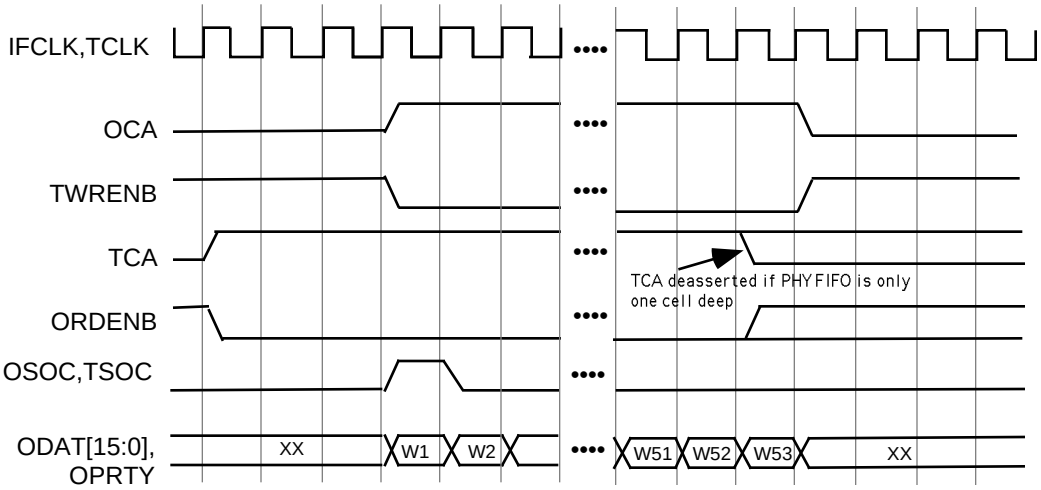

The ORDENB signal at the RCMP is generated from the inverted TCA signal from the PHY device. The TWRENB signal at the PHY device is generated from the inverted OCA signal from the RCMP. This scheme works as follows: TCA is asserted when the PHY device has space for at least one cell. This causes ORDENB to be asserted which signals the RCMP to output a cell. However, the RCMP only outputs a cell if the cell is available, in which case OCA is asserted. At the same time, TWRENB is also asserted, and the first word of the cell is sent out. At the end of the cell transfer, OCA is deasserted, causing TWRENB to be deasserted also. *It is important to ensure that OCA is only deasserted at the end of the cell transfer*, not 4 words before the end. There is a

programmable option for OCA deassertion in the RCMP (Register 38, Bit 1). This bit is set to 1 by default, indicating OCA deassertion at the end of cell transfer.

Note that in the scenario just described, we have assumed that there is always space available on the PHY device, and the actual transfer is determined by the cell availability at the RCMP. The other scenario is when the space availability on the PHY device is the determining factor. In this scenario, OCA is asserted with a cell waiting to be sent out by the RCMP, and TWRENB is also asserted. No cell transfer takes place until the PHY device has cell space, and TCA gets asserted, which causes ORDENB to be asserted. With TWRENB asserted but no space in the PHY device to receive a cell, the PHY device will likely indicate an input FIFO over-run condition. Note that the PHY device should not allow data to be written in when TCA is deasserted. In general, this scenario will not occur if the PHY device has enough bandwidth for the cell traffic from the RCMP. The first scenario in the previous paragraph is a more likely case, since the RCMP can be dropping cells due to policing, thus reducing the rate of cells being sent out. In this case, TCA will always be asserted, indicating that the PHY device always has space.

The RCMP can be programmed to output an inverted OCA, thus avoiding the use of the inverter. Refer to Register 38, Bit 14, which is set to 0 by default, indicating no inversion on OCA.

SUMMARY

An overview has been given for the use of the RCMP in the egress direction to provide address translation, OAM processing, policing and moitoring functions.

In the egress application of the RCMP, two cases of the ATM cell interface have been described. Synchronous FIFO's are used to buffer the master-master interface between the switch core and the RCMP. Only inverters are needed for the slave-slave interface between the RCMP and the PHY device. Note that the principles of the interconnection schemes described can be applied in general to any master-master and slave-slave SCI-PHY interfaces.

APPENDIX 1 EGRESS LOOP APPLICATION

Introduction

In a typical switch port card, the RCMP is used only in the Ingress direction, and a separate device is needed to implement the Egress function. However, for cases where the only Egress functions needed are address translation (which in most cases just means stripping off the appended words used by switch fabric) and multicasting, then a separate Egress device can be avoided. That is, the Ingress RCMP can be used to provide these two Egress functions, provided that there is spare bandwidth available in the RCMP. This spare bandwidth exists in most PDH applications where the RCMP is used in the multi-PHY mode interfacing with multiple T1/E1 PHY ports. For example, if 32 E1 PHY ports are interfaced with the RCMP-200 (PM7323), then the total Ingress bandwidth will only be 65.5 Mbps, which leaves a lot of spare bandwidth since the RCMP-200 can handle 200 Mbps.

Fig. A1 shows the system configuration that implements this function. This is referred to as the Egress Loop application. The cell traffic coming from the switch fabric are looped back into the input cell interface of the RCMP, which treats these cells as if they are sourced from another PHY port. The RCMP process these cells by doing the VC Table lookup, and performs the address translation and multicasting if appropriate. At the output of the RCMP, an external piece of glue logic routes the Egress cell traffic to the Multi-PHY interface glue logic, which directs the traffic to the appropriate PHY port.

Fig. A1 Egress Loop System Configuration

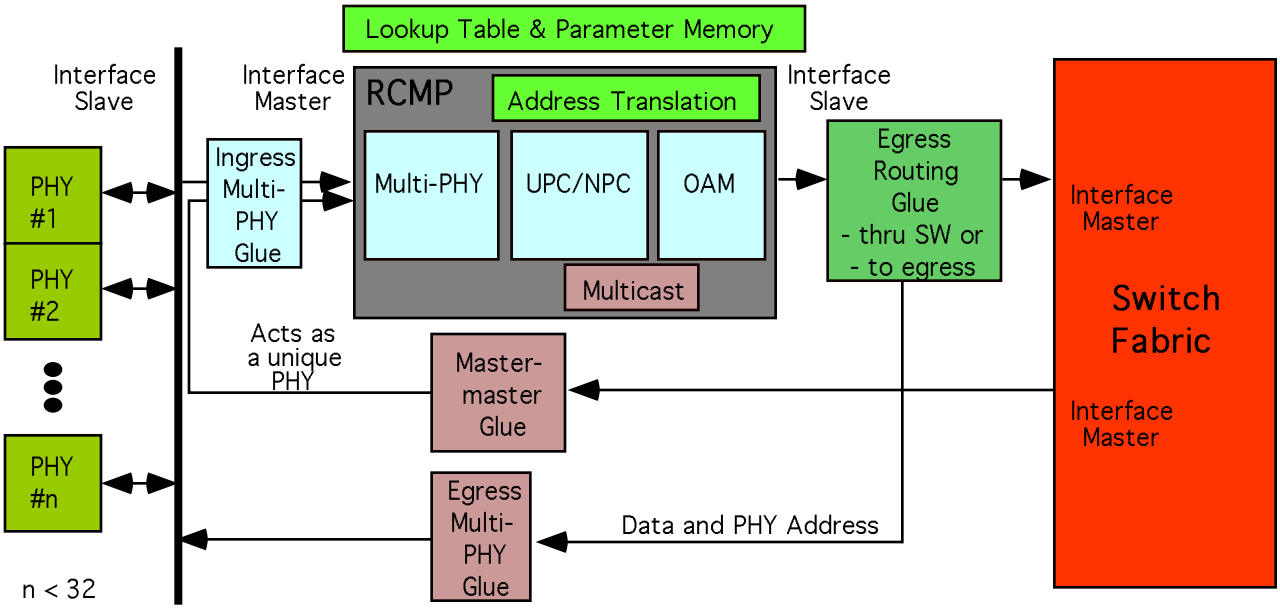

The following provides a detailed description of each of the functional blocks that are required to provide this Egress Loop function.

Ingress Multi-PHY Glue Logic

In the Multi-PHY application, the data bus is shared between the multiple PHY ports, and the RCMP arbitrates between the PHY ports by address polling. Normally, a maximum of 32 PHY ports can be supported. However, since the Egress cell traffic is presented as an Ingress PHY port, only up to 31 PHY ports are supported.

This glue logic block performs PHY address decoding and cell-available signal selection. The RCMP outputs 5 bits of PHY address, which needs to be decoded to generate the read-enable signal to the appropriate PHY port. One out of up to 31 cell-available signals are selected based on the decoded PHY address. Fig. A2 shows the glue logic required. The Egress traffic is represented by PHY "port" #32.

Fig. A2 Ingress Multi-PHY Glue Logic

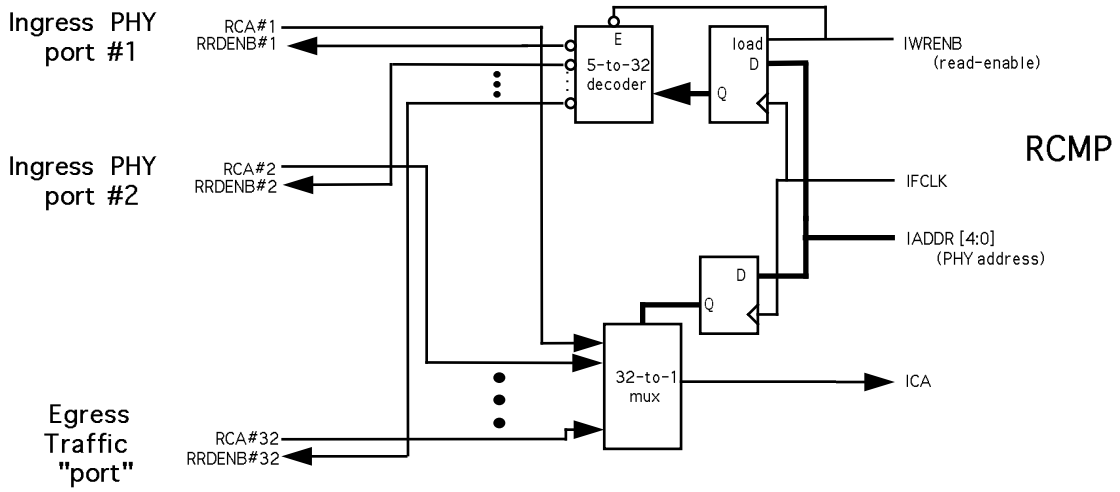

Egress Multi-PHY Glue Logic

This block performs a very similar function as the Ingress Multi-PHY glue logic. It decodes the PHY address to generate the write-enable signal to the appropriate PHY port. It selects one out of 31 cell-available signals based on the decoded PHY address.

Master-Master Glue Logic

The Egress cell traffic is looped back into the input of the RCMP as if it is from another Ingress PHY port. The input of the RCMP acts as the interface master. If the switch

output is an interface slave, then the interface is straight-forward and no extra glue logic is required. However, if the switch is an interface master, which is the case in general, then the master-master glue logic is needed, which is described in the "Master-Master Interface" section on page 2.

Egress Routing Glue Logic

The function of this block is to route the cell traffic from the output of the RCMP to either the switch fabric or one of the 31 Egress PHY ports, and to provide the necessary control signals. This is done by using the first prepend word[8] of a cell to indicate the intended destination of the cell. This prepend word contains a routing bit that indicates whether the cell should be routed to the switch or back to the Egress direction. It also contains a 5-bit PHY address for the Egress cells. The reason for using the first prepend word is to allow the glue logic to determine the routing of the cell before the remainder of the cell (which includes possible pre/postpend words to be used by the switch) arrives, thus avoiding the need to buffer the cell.

Note that there are three *assumptions* made on the system for Egress routing:

1. The same clock is used for the PHY interface and the Switch interface (ie. the input and the output interfaces of the RCMP). This is because the Egress Routing glue logic does not perform synchronization between two different clock domains. On the other hand, the Master-master glue logic, being a FIFO implementation, can operate with different clocks the two sides.

2. Both the PHY interface and the switch inteface have the same data bus width (8 or 16 bit). This greatly simplifies the Egress Routing glue logic. Otherwise, the glue logic will have to convert between a 16-bit and an 8-bit bus, which adds memory and more complicated control logic.

3. Cells coming out of the switch fabric do not have prepend/postpend words, since they have to match the cell length of those coming from the PHY ports. Otherwise, the master-master glue logic will have to be able to strip off these prepend/postpend words.

The glue logic should perform the following steps for every cell transfer:

1. Glue logic asserts ORDENB to the RCMP to read out the first prepend word, and then halts the cell read by deasserting ORDENB.

2. Examine the routing bit to connect the control signals (read-enable and cell-available signals) from either the switch or the Egress Multi-PHY glue logic.

3. Strip off all prepend words for Egress cells; otherwise go to Step 4. This is because the RCMP has a global configuration that sets the number of pre/postpend words, which

[8]The RCMP has the flexibility to be configured to output up to 5 pre/postend words for a 16-bit interface, or up to 10 pre/postpend bytes for an 8-bit interface. Note that the term "word" is used here to refer to both 8 or 16-bit data.

affects both Ingress and Egress cells. Since switches in general do require the pre/postpend words for routing, these pre/postpend words have to be stripped off so that only the ATM-cell (53-bytes long) is presented to the PHY ports. Prepend words are stripped off by asserting the ORDENB signal to empty out these words from the RCMP output FIFO but not actually transferring the data to the destination.

4. Start cell transfer. If the cell is destined to the switch, the switch controls the ORDENB signal. Otherwise, the glue logic controls both the ORDENB and TWRENB (to the PHY device) signals.

5. Strip off the postpend words for Egress cells; otherwise continue the transfer until the end of postpend words. Postpend words are stripped off by asserting the ORDENB signal to empty out these words from the RCMP output FIFO but not actually transferring the data to the destination. From the Egress PHY port point of view, the cell transfer already ended before the postpend stripping starts.

For the Egress cell flow, since both the RCMP and the PHY ports have slave interfaces, the Egress Routing glue logic acts as the master, providing both the ORDENB and TWRENB signals.

Egress Cell Transfer

Fig. A3 Egress Routing Glue Logic Timing Diagram

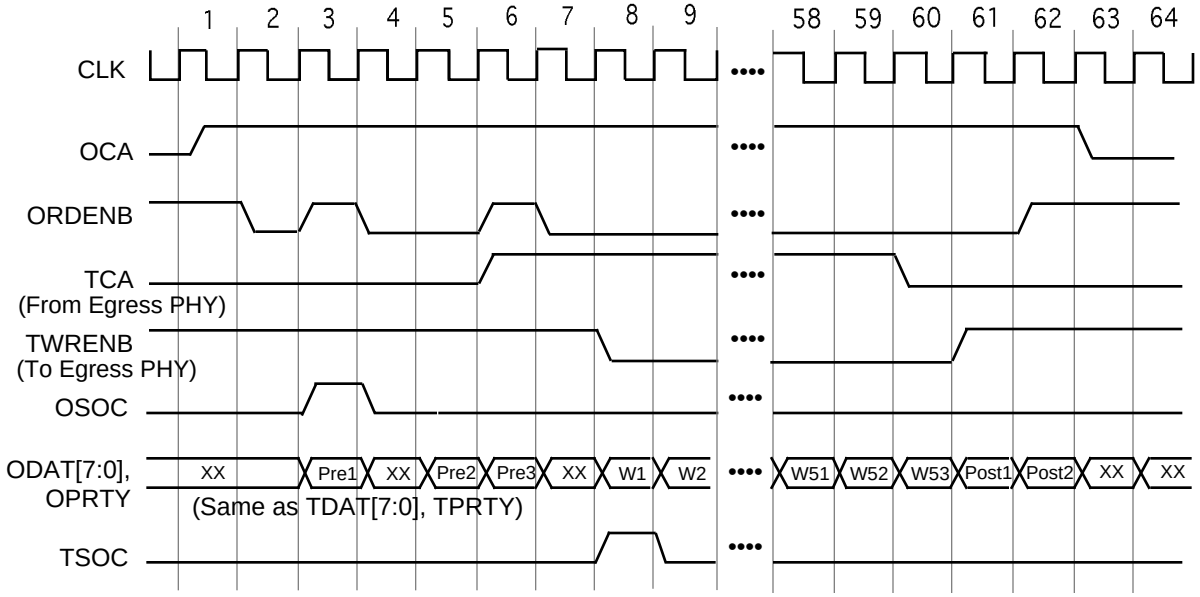

Fig. A3 shows the timing diagram of a typical Egress cell transfer for an 8-bit interface. The following describes the sequence of events for the transfer of a cell with 3 prepend words (the first one for Egress routing) and 2 postpend words.

Clock cycle 1: RCMP indicates a cell is available by asserting OCA.

Clock cycle 2: Glue logic responds by asserting ORDENB to read out the first prepend word.

Clock cycle 3: The first prepend word along with OSOC are read out from the RCMP. Glue logic deasserts ORDENB to allow the routing information to be parsed in cycle 4.

Clock cycle 4: Once the glue logic has determined that the cell is headed towards the Egress direction, it asserts ORDENB to read the second prepend word from the RCMP. Note that even though these prepend words are on the data bus, they are ignored by the PHY ports, since TWRENB is not asserted.

Clock cycle 5: Second prepend word is read from the RCMP. Note that ORDENB is still held asserted to read the last prepend word from the RCMP.

Clock cycle 6: Third prepend word is read from the RCMP. ORDENB deasserted while waiting for the PHY port to assert TCA[9].

Clock cycle 7: TCA is sampled high by the RCMP, and ORDENB is therefore asserted to start the cell transfer.

Clock cycle 8: The first word of the ATM cell is read from the RCMP. TWRENB is asserted by the glue logic to align with the data. Here, TSOC is generated by the glue logic to align with the first word.

Clock cycle 9-60: The rest of the ATM cell is transferred. In clock cycle 60, TCA deasserted by the PHY port.

Clock cycle 61: TWRENB is deasserted by the glue logic to end the cell transfer to the PHY port. Glue logic still holds ORDENB asserted to read out the postpend words from the RCMP.

Clock cycle 62: ORDENB is deasserted by the glue logic as the last postpend word of the ATM is transferred.

Clock cycle 63: OCA deasserted by RCMP. End of cell transfer.

Note that one consequence of this Egress Routing glue logic is that the first prepend word cannot be used by the switch fabric. The maximum number of pre/postpend words thus becomes 9.

[9]This timing diagram shows the case where TCA from the PHY port determines when the actual cell transfer starts. However, for the case where there is space in the PHY port all the time, TCA is constantly asserted, and the glue logic will not assert TWRENB until all the prepend words have been read out from the RCMP. (eg. the earliest clock cycle where TWRENB is asserted is cycle 7, such that Word1 will immediately follow Prepend Word3).

Fig. A4 shows the state diagram of the Egress Routing Glue logic and Fig. A5 shows the block diagram.

Fig. A4 State Diagram of the Egress Routing Glue Logic

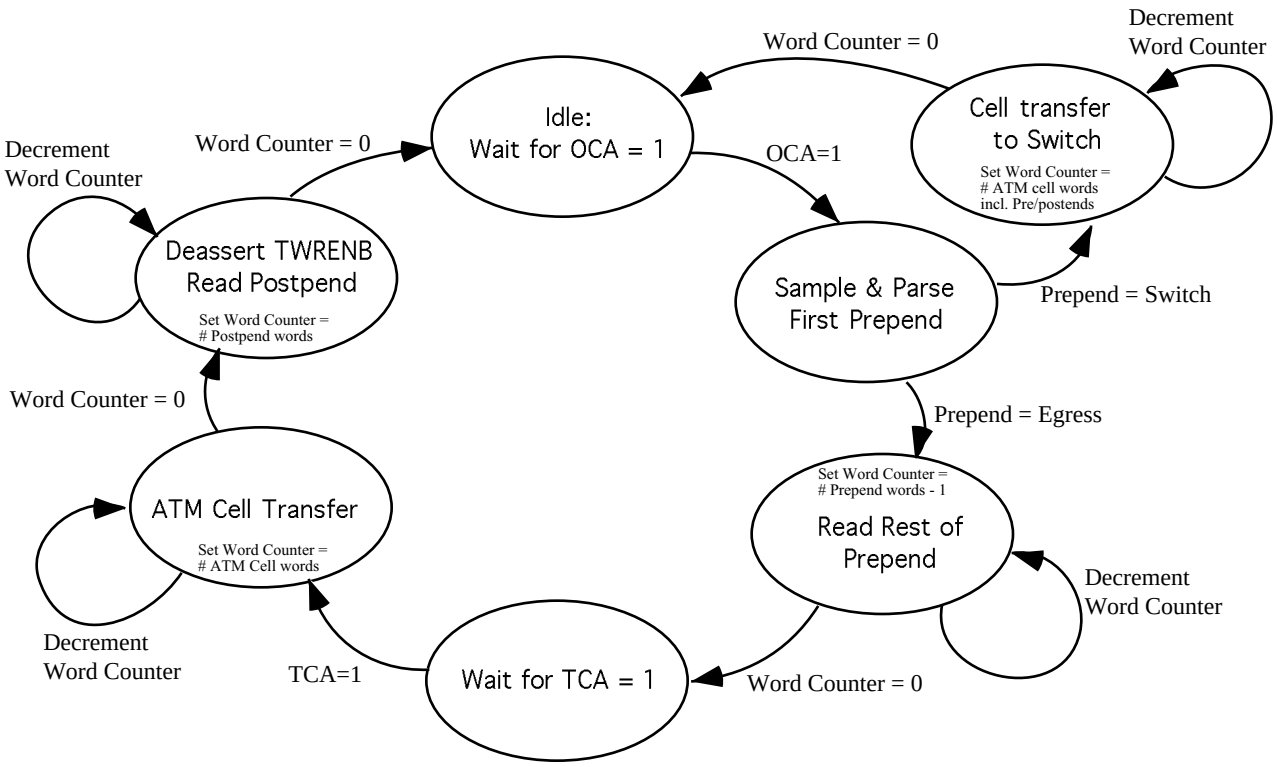

Fig. A5 Block Diagram of the Egress Routing Glue Logic

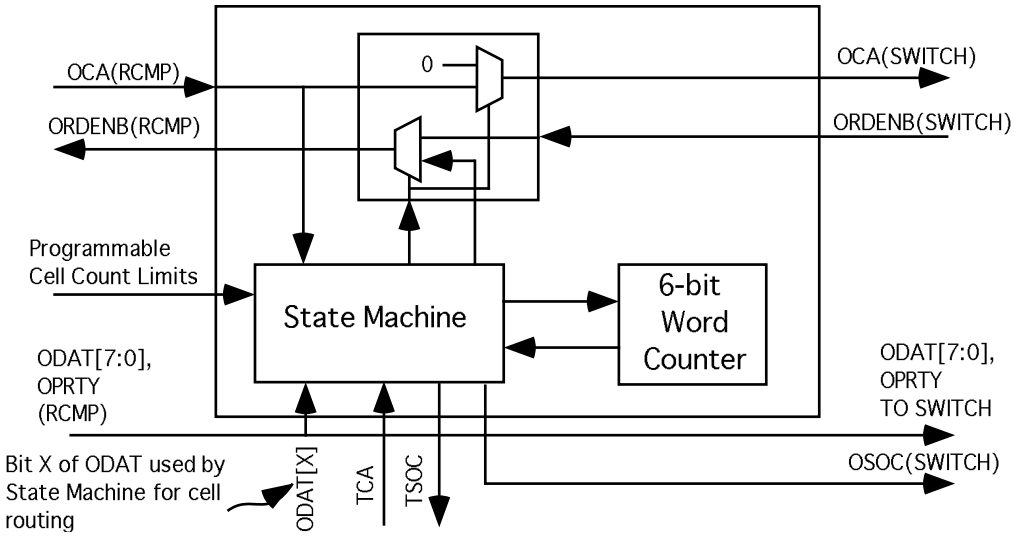

Summary

The Egress Loop application of the RCMP has been described in the above. By using the spare bandwidth of the RCMP in the Ingress direction for Egress cell processing, a separate IC for the Egress function is not required at all. In fact, all the glue logic described can be implemented in either PLD's or PAL's, thus completely avoiding the need to develop an ASIC for the switch port card.

NOTES

Contact us for applications support:

FAX: (604) 668-7301

PHONE: (604) 668-7300

Email: apps@pmc-sierra.bc.ca

PMC-951205P2 Issue date: March 13, 1996